

A Study of Unrestricted Abstract Argumentation Frameworks

Ringo Baumann

Leipzig University, Germany
baumann@informatik.uni-leipzig.de

Christof Spanring

University of Liverpool, United Kingdom
TU Wien, Austria
christof.spanring@tuwien.ac.at

Abstract

Research in abstract argumentation typically pertains to finite argumentation frameworks (AFs). Actual or potential infinite AFs frequently occur if they are used for the purpose of nonmonotonic entailment, so-called instantiation-based argumentation, or if they are involved as modeling tool for dialogues, n-person-games or action sequences. Apart from these practical cases a profound analysis yields a better understanding of how the nonmonotonic theory of abstract argumentation works in general. In this paper we study a bunch of abstract properties like SCC-recursiveness, expressiveness or intertranslatability for unrestricted AFs.

1 Introduction

In the last two decades formal argumentation has become one major research area in Artificial Intelligence. The main reason for its popularity is its numerous fields of application ranging from nonmonotonic reasoning over multi-agent systems as well as aspects of knowledge representation to analysis tool for debates or dialogues (cf. [Bench-Capon and Dunne, 2007; Rahwan and Simari, 2009] for excellent overviews). Dung's abstract argumentation frameworks (AFs) [Dung, 1995], which are essentially directed graphs, play a dominant role in this field since they are frequently used as an evaluation component in deductive approaches to argumentation and as common basis for more complex abstract argumentation formalisms. It is therefore not surprising that many fundamental properties typically considered for nonmonotonic formalisms have been studied in recent years. Among others, central issues like replaceability [Oikarinen and Woltran, 2011], expressibility [Dunne *et al.*, 2015] intertranslatability [Dvořák and Woltran, 2011] as well as general modularity and locality properties were studied [Baroni and Giacomin, 2007; Baumann, 2011; Baroni *et al.*, 2014]. As a matter of fact, almost all existing research on abstract argumentation is restricted to finite AFs, and thus there is no systematic analysis of abstract properties in the unrestricted case. Studying infinite AFs is far from being an academic exercise since it is motivated by the following observations/facts:

- *actual infinity*: The semantics of rule-based argumentation formalisms (like ASPIC [Prakken, 2010]) is usually

given via the evaluation of induced Dung-style AFs. It was already observed that even a finite set of rules may lead to an infinite set of arguments (cf. [Caminada and Oren, 2014; Strass, 2015]). Consequently, after clarifying whether some argumentation semantics always provides one with a formal meaning one may ask for sound and complete procedures evaluating even infinite AFs.

- *potential infinity*: Allowing an infinite number of arguments is essential in applications where upper bounds on the number of available arguments cannot be established a priori. For instance, this applies to multi-agent dialogues [Belardinelli *et al.*, 2015] as well as modeling approaches including time or action sequences [Baumann and Strass, 2012]. A central issue regarding dynamics is whether, and if so how, one may simplify a current framework in a semantically neutral way even if an unbounded number of arguments may arise in future.
- *theoretical foundations*: The first theoretical results as well as applications regarding infinite AFs can be traced back to Dung himself [Dung, 1995]. For instance, he showed how to model (possibly infinitely many) solutions of n-person games with the help of AFs. Surprisingly, this research direction was somehow neglected for nearly 20 years. In recent times infinite AFs are gaining attention [Baroni *et al.*, 2013; Baumann and Spanring, 2015]. Besides the deeper understanding of how the nonmonotonic theory of abstract argumentation works, we believe that a profound analysis will be of great benefit for the argumentation community as well as other research fields using Dung-style AFs as modeling tool.

The paper is organized as follows. Sect. 2 provides the necessary background on AFs. Sect. 3 contains our main contributions, namely an analysis of unrestricted AFs w.r.t. computational aspects including the characteristic function as well as SCC-recursiveness, subset relations between semantics, replaceability, expressibility and intertranslatability. Moreover, for each considered topic we contrast our general results with former ones restricted to finite AFs. Due to the limited space we omit all proofs. Nevertheless, in almost all cases we included some short comments indicating how to prove the statement in question. Finally, in Sect. 4 we discuss our results and give some pointers for future work.

Table 1: Existence and Uniqueness

	<i>stb</i>	<i>ss</i>	<i>stg</i>	<i>cf2</i>	<i>stg2</i>	<i>pr</i>	<i>ad</i>	<i>co</i>	<i>na</i>	<i>cf</i>	<i>gr</i>	<i>il</i>	<i>eg</i>
finite AFs		∃	∃	∃	∃	∃	∃	∃	∃	∃	∃!	∃!	∃!
unrestricted AFs						∃	∃	∃	∃	∃	∃!	∃!	∃

2 Background

More often than not the assumption of a certain reference set \mathcal{U} , say *universe of arguments* is implicitly assumed and not explicitly stated in argumentation papers. For any AF $F = (A, R)$ we require that $A \subseteq \mathcal{U}$ as well as $R \subseteq A \times A$. In order to deal with AFs possessing infinitely many arguments we have to request that $|\mathcal{U}| \geq \aleph_0 = |\mathbb{N}|$. No further conditions are imposed. An AF $F = (A, R)$ is called *finite* if $|A| \in \mathbb{N}$, otherwise it is called *infinite*. Moreover, an AF F is called *finitary* if each of its arguments has only finitely many attackers. We use \mathcal{F} as an abbreviation for the class of *unrestricted frameworks*, i.e. \mathcal{F} contains all AFs w.r.t. \mathcal{U} .

We say a attacks b in $F = (A, R)$ whenever $(a, b) \in R$. An argument $a \in A$ is *defended* by a set $A' \subseteq A$ in F if for each $b \in A$ with $(b, a) \in R$, b is attacked by some $a' \in A'$ in F . For a set $E \subseteq A$ we use E_F^+ or simply, E^+ for $E \cup \{b \mid (a, b) \in R, a \in E\}$. If $G = (B, S)$, we use $A(G)$ as well as $R(G)$ to refer to the first or second component of G , i.e. B or S , respectively. The set $L(G) = \{a \in A(G) \mid (a, a) \in R(G)\}$ contains all self-attacking arguments. Finally, the union of two AFs as well as the restriction to a certain set $C \subseteq \mathcal{U}$ is defined as expected, i.e. $F \cup G = (A(F) \cup A(G), R(F) \cup R(G))$ and $F|_C = (A(F) \cap C, R(F) \cap (C \times C))$.

An *extension-based semantics* $\sigma : \mathcal{F} \rightarrow 2^{2^{\mathcal{U}}}$ is a function which assigns to any AF F a set of sets of arguments $\sigma(F) \subseteq 2^{A(F)}$. Each one of them, a so-called σ -*extension* E , is considered to be acceptable with respect to F . For two semantics σ and τ we write $\sigma \subseteq \tau$, if for any AF F , $\sigma(F) \subseteq \tau(F)$. Besides conflict-free and admissible sets (abbreviated by *cf* and *ad*) we consider a large number of mature semantics, namely naive, stage, stable, semi-stable, complete, preferred, grounded, ideal and eager semantics (abbreviated by *na*, *stg*, *stb*, *ss*, *co*, *pr*, *gr*, *il* and *eg* respectively). A recent overview can be found in [Baroni *et al.*, 2011].

Definition 1 Let $F = (A, R)$ be an AF and $E \subseteq A$.

1. $E \in cf(F)$ iff for no $a, b \in E$, $(a, b) \in R$,
2. $E \in na(F)$ iff $E \in cf(F)$ and for no $I \in cf(F)$, $E \subset I$,
3. $E \in stg(F)$ iff $E \in cf(F)$ and for no $I \in cf(F)$, $E^+ \subset I^+$
4. $E \in stb(F)$ iff $E \in cf(F)$ and $E^+ = A$,
5. $E \in ad(F)$ iff $E \in cf(F)$ and E defends all its elements,
6. $E \in ss(F)$ iff $E \in ad(F)$ and for no $I \in ad(F)$, $E^+ \subset I^+$
7. $E \in co(F)$ iff $E \in ad(F)$ and for any $a \in A$ defended by E in F , $a \in E$,
8. $E \in pr(F)$ iff $E \in co(F)$ and for no $I \in co(F)$, $E \subset I$,
9. $E \in gr(F)$ iff $E \in co(F)$ and for no $I \in co(F)$, $I \subset E$,
10. $E \in \sigma(F)$ iff $E \in ad(F)$, $E \subseteq \bigcap \tau(F)$ (for $\sigma/\tau \in \{il/pr, eg/ss\}$) and there is no $I \in ad(F)$ satisfying $E \subset I \subseteq \bigcap \tau(F)$. \diamond

3 From Finite to Unrestricted Frameworks

A priori it is unclear whether certain abstract properties satisfied for finite AFs carry over to unrestricted ones. To the best of our knowledge, the only properties which were already analyzed outside the realm of finiteness are *universal* and *unique definedness* of semantics (cf. [Baumann and Spanring, 2015] for an overview). These results vividly show that it makes an important difference whether we restrict ourselves to finite AFs. After a brief review of these results we will tackle a bunch of well-known and important issues in the theory of abstract argumentation.

3.1 Existence and Uniqueness

Already Dung considered the question of *existence* and *uniqueness* of extensions for unrestricted frameworks. For instance, he showed that preferred semantics guarantees at least one extension for any AF [Dung, 1995, Corollary 12]. Later on further results were achieved showing that existence and uniqueness of extensions may depend on the structure of the considered AFs [Weydert, 2011; Baumann and Spanring, 2015]. Table 1 provides a compact overview. An empty entry indicates a possible collapse of the semantics σ , i.e. for some AFs F we observe $\sigma(F) = \emptyset$. The entries "∃" or "∃!" indicate that for any AF F , $|\sigma(F)| \geq 1$ or $|\sigma(F)| = 1$, respectively.

3.2 Computing the Grounded Extension

In contrast to Definition 1 Dung defined the grounded extension of an AF F as the \subseteq -least fixpoint of the so-called *characteristic function* $\Gamma_F : 2^{A(F)} \rightarrow 2^{A(F)}$ with $E \mapsto \{a \in A(F) \mid a \text{ is defended by } E\}$. Moreover, he showed that this definition coincides with the \subseteq -least complete extension [Dung, 1995, Theorem 25]. In case of finite (and even finitary) AFs the characteristic function is ω -continuous, implying that the least fixpoint equals $\bigcup_{i < \omega} \Gamma_F^i(\emptyset)$ [Gabbay and Grossi, 2014]. However, in case of non-finitary AFs Γ_F might not be ω -continuous but still remains monotonic. Consequently, ω -iteration might not be sufficient (cf. Example 1). Nevertheless, even in case of unrestricted AFs we may ensure the stepwise "computation" of the grounded extension by using ordinals beyond ω . The assertion can be shown with results from lattice theory.

Theorem 1 For any *unrestricted AF* $F = (A, R)$ exists an ordinal α of size at most $|A(F)|$, s.t. $gr(F) = \{\bigcup_{i < \alpha} \Gamma_F^i(\emptyset)\}$.

Example 1 ($\omega \cdot 2$ -iteration) Consider the following non-finitary AF $F = (A, R)$ where $A = \{a_i, b_i \mid i \in \mathbb{N}\}$ and $R = \{(a_i, a_{i+1}), (b_i, b_{i+1}), (a_{2i}, b_1) \mid i \in \mathbb{N}\}$.




Table 2: SCC-recursiveness and Well-definedness

	<i>stb</i>	<i>ss</i>	<i>stg</i>	<i>cf2</i>	<i>stg2</i>	<i>pr</i>	<i>ad</i>	<i>co</i>	<i>na</i>	<i>cf</i>	<i>gr</i>	<i>il</i>	<i>eg</i>
finite AFs	SW	W	W	SW	SW	SW	SW	SW	W	SW	SW	W	W
unrestricted AFs	W	W	W	S	S	W	W	W	W	W	W	W	W


Clearly, $gr(F) = \{\{a_i, b_i \mid i \in \mathbb{N}, i \text{ is odd}\}\}$. Let us compute some iterations. For finite ordinals n we have $\bigcup_{i=1}^n \Gamma_F^i(\emptyset) = \{a_1, \dots, a_{2n-1}\}$. Then, the first limit ordinal ω yields $\bigcup_{i < \omega} \Gamma_F^i(\emptyset) = \{a_i \mid i \in \mathbb{N}, i \text{ is odd}\}$ which does not coincide with the grounded extension. Pursuing with successor ordinals results in $\bigcup_{i=1}^{\omega+n} \Gamma_F^i(\emptyset) = \{a_i \mid i \in \mathbb{N}, i \text{ is odd}\} \cup \{b_1, \dots, b_{2n-1}\}$ and finally, after $\omega + \omega$, i.e. $\omega \cdot 2$ -iterations we end up with a fixpoint coinciding with the grounded extension. Please observe that the cardinality of $\omega \cdot 2$ equals $\aleph_0 = |A(F)|$ as promised by Theorem 1.

3.3 SCC-recursiveness

The concept of *SCC-recursiveness* plays a decisive role in the study of argumentation semantics. We mention two important facts: firstly, it guarantees desirable properties like *directionality* and secondly, it already gave rise for new semantics definitions like *cf2* and *stg2* which are able to overcome certain questionable results of former semantics (cf. [Baroni *et al.*, 2005; Baroni and Giacomin, 2007; Dvořák and Gaggl, 2012] for more details). Semantics satisfying *SCC-recursiveness* can be incrementally computed via a recursive schema which is based on the decomposition of AFs along their *strongly connected components (SCCs)*. More precisely, the schema takes a *base semantics* σ , proceeds along the induced partial ordering and evaluates the SCCs according to σ while propagating relevant results to subsequent SCCs. This procedure defines a $\sigma 2$ -semantics. We say that a semantics τ is *SCC-recursive* if there is a base function σ , s.t. for any AF F , $\tau(F) = \sigma 2(F)$. In case of finite AFs a bunch of semantics have been identified to be SCC-recursive (cf. [Baroni *et al.*, 2005; Gaggl and Strass, 2014], Table 2 gives an overview).

In case of unrestricted frameworks we are faced with some difficulties in drawing conclusions. Firstly, infinite AFs need not possess initial SCCs which is granted for finite AFs. This makes checking whether a certain set is a $\sigma 2$ -extension more complicated and due to the recursive definitions not that easy to handle. Secondly, even worse, even if initial SCCs are available there is no guarantee that the recursion will stop in finitely many steps.

Example 2 (Non-SCC-recursiveness) Consider the following finitary AF $F = (A, R)$ where $A = \{a_i, b_i \mid i \in \mathbb{N}\}$ and $R = \{(b_i, a_i), (a_{i+1}, a_i), (a_i, b_{i+1}) \mid i \in \mathbb{N}\}$.



Obviously, $B = \{b_i \mid i \in \mathbb{N}\} \in \tau(F)$ for any semantics $\tau \in \{stb, ss, stg, pr, ad, co, na, cf, gr, il, eg\}$. Striving for a contradiction, suppose that there is base semantics σ , s.t.

$\tau(F) = \sigma 2(F)$. Observe that the AF F possesses two SCCs, namely $C_1 = \{b_1\}$ and $C_2 = A \setminus \{b_1\}$. For C_1 we end up with the base case returning $\{b_1\} \in \sigma(F|_{C_1})$. For C_2 we have to consider the AF $F' = F|_{C_2 \setminus \{a_1\}}$ (since a_1 is attacked by the SCC C_1). Now, we have to determine whether $B' = \{b_i \mid i \in \mathbb{N}, i \geq 2\}$ is a $\sigma 2$ -extension of F' . Note that this is equivalent to deciding whether B is a $\sigma 2$ -extension of F . This means, the candidate set B leads to infinite recursion, i.e. the base case will never be considered although $B \in \tau(F)$. \diamond

In contrast to all other semantics available in the literature, *cf2* and *stg2* were originally defined in terms of the recursive schema with naive and stage semantics as base functions. This means, SCC-recursiveness is given by definition but Example 2 shows that both semantics are not well-defined (in the sense of a function) in the general unrestricted case since we cannot determine whether $B \in cf2(F)$ or $B \in stg2(F)$, respectively.

Theorem 2 For *unrestricted frameworks* we have,

- for any $\sigma \in \{stb, ss, stg, pr, ad, co, na, cf, gr, il, eg\}$, semantics σ is not SCC-recursive and
- cf2* and *stg2* are SCC-recursive.

Table 2 summarizes all results. The entries "SW", "S" and "W" in row *certain* AFs and column σ mean that for *certain* AFs the semantics σ is SCC-recursive and well-defined, SCC-recursive but not well-defined and well-defined but not SCC-recursive.

3.4 Subset Relations between Semantics

Note that the shorthand $\sigma \subseteq \tau$ requires that both semantics are total functions since a framework to which one of these semantics is undefined renders the subset shorthand undefined itself. Consequently, $stb \subseteq stg2 \subseteq cf2 \subseteq na$ does not hold in general (Example 2) but it does in case of finite AFs as shown in [Dvořák and Gaggl, 2012]. All other subset relations already known from finite AFs carry over to unrestricted AFs which can be shown with reasonable effort. The following figure summarizes the results in a compact way. For two semantics σ and τ , $\sigma \subseteq \tau$ iff there is a path from σ to τ in Figure 1. A red-highlighted arrow indicates that the corresponding subset relation is guaranteed for finite frameworks only.

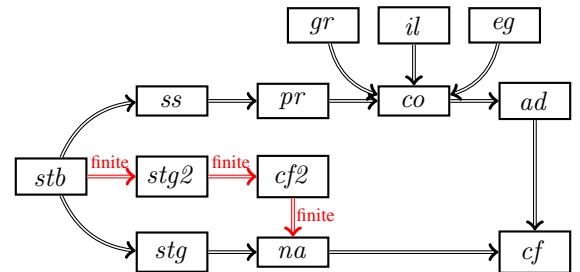


Figure 1: Subset Relations

Table 3: Expansion Equivalence

	<i>stb</i>	<i>stg</i>	<i>ss</i>	<i>pr</i>	<i>ad</i>	<i>il</i>	<i>eg</i>	<i>co</i>	<i>gr</i>	<i>na</i>	<i>cf</i>
finite (and jointly exp.) AFs	$k(stb)$	$k(stb)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(co)$	$k(gr)$	$k(na)$	$k(na)$
unrestricted AFs	?	?	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$?	?	$k(na)$	$k(na)$

3.5 Expansion Equivalence

In contrast to classical (monotone) logics *ordinary equivalence* is not sufficient to capture replaceability without loss of information (cf. [Baumann and Strass, 2016] for more detailed information on this issue). This means, analogously to other non-monotonic logics one may easily find two AFs F and G which possess the same σ -extensions but differ semantically if augmented by a further AF H . We say that both frameworks are *expansion* (or *strong*) *equivalent* if the latter is impossible.

Definition 2 Given a semantics σ . Two AFs F and G are

1. *ordinarily equivalent* w.r.t. σ iff $\sigma(F) = \sigma(G)$,
2. *expansion equivalent* w.r.t. σ (for short, $F \equiv_E^\sigma G$) iff for each AF H we have, $\sigma(F \cup H) = \sigma(G \cup H)$, \diamond

It was the main result in [Oikarinen and Woltran, 2011] that expansion equivalence can be decided by looking at the syntax only. More precisely, they introduced the notion of a *kernel* of an AF F which is (informally speaking) a graph where *redundant* attacks w.r.t. F are deleted or added and showed that syntactical identity of suitably chosen kernels characterizes expansion equivalence w.r.t. the considered semantics.

Definition 3 Let $\sigma \in \{stb, ad, gr, co\}$. For any AF $F = (A, R)$ we define the σ -*kernel* $F^{k(\sigma)} = (A, R^{k(\sigma)})$ of F as:

1. $R^{k(stb)} = R \setminus \{(a, b) \mid a \neq b, (a, a) \in R\}$,
2. $R^{k(ad)} = R \setminus \{(a, b) \mid a \neq b, (a, a) \in R, \{(b, a), (b, b)\} \cap R \neq \emptyset\}$,
3. $R^{k(gr)} = R \setminus \{(a, b) \mid a \neq b, (b, b) \in R, \{(a, a), (b, a)\} \cap R \neq \emptyset\}$,
4. $R^{k(co)} = R \setminus \{(a, b) \mid a \neq b, (a, a), (b, b) \in R\}$,
5. $R^{k(na)} = R \cup \{(a, b) \mid a \neq b, \{(a, a), (b, a), (b, b)\} \cap R \neq \emptyset\}$. \diamond

Please note that all introduced kernels k are node- and loop-preserving. This means, for any AF F , $A(F) = A(F^k)$ and $L(F) = L(F^k)$. We say that k *characterizes an equivalence notion* \equiv if for any two AFs F and G , $F \equiv G$ iff $F^k = G^k$.

Theorem 3 [Oikarinen and Woltran, 2011; Baumann et al., 2016b; Gaggl and Woltran, 2013] *For finite AFs, we have:*

1. $k(stb)$ characterizes \equiv_E^σ for any $\sigma \in \{stb, stg\}$,
2. $k(ad)$ characterizes \equiv_E^σ for any $\sigma \in \{ss, eg, ad, pr, il\}$,
3. $k(\sigma)$ characterizes \equiv_E^σ for any $\sigma \in \{gr, co\}$ and
4. $k(na)$ characterizes \equiv_E^σ for any $\sigma \in \{na, cf\}$.

We mention that all semantics considered in the literature satisfy *regularity* given that finite AFs are considered. Regularity means that expansion equivalent AFs have to share the same arguments [Baumann, 2012]. This is obvious for the semantics considered in Theorem 3 since all characterizing kernels are node-preserving. Moreover, since the background universe \mathcal{U} contains infinitely many arguments it is always possible to expand finite AFs F and G by "fresh" arguments, i.e. arguments which do not occur in $A(F) \cup A(G)$. Clearly, dropping the finite assumption may rule out this possibility.

Example 3 (Non-regularity) Given $c \in \mathcal{U}$, $F = (\mathcal{U} \setminus \{c\}, \{(a, a) \mid a \in \mathcal{U} \setminus \{c\}\})$ and $G = (\mathcal{U}, \{(a, a) \mid a \in \mathcal{U} \setminus \{c\}\})$. For any H we observe $stb(F \cup H) = stb(G \cup H)$. In particular,

$$stb(F \cup H) = \begin{cases} \{\{c\}\}, & \text{if } \{(c, a) \mid a \in \mathcal{U} \setminus \{c\}\} \subseteq R(H) \\ & \text{and } (c, c) \notin R(H) \\ \emptyset, & \text{otherwise} \end{cases}$$

Consequently, $F \equiv_E^{stb} G$ although $A(F) \neq A(G)$. \diamond

Example 4 (Non-coincidence) Let $F = (\mathcal{U}, \{(a, a) \mid a \in \mathcal{U}\})$ and $G = (\mathcal{U}, \{(a, b) \mid a, b \in \mathcal{U}, a \neq b\})$. Applying the grounded kernel does not change anything for either framework, i.e. $F^{k(gr)} = F$ and $G = G^{k(gr)}$. Due to the absence of unattacked arguments we deduce $\Gamma_F(\emptyset) = \Gamma_G(\emptyset) = \emptyset$ implying that $gr(F \cup H) = gr(G \cup H) = \{\emptyset\}$ for any AF H . Consequently, $F \equiv_E^{gr} G$ although $F^{k(gr)} \neq G^{k(gr)}$. \diamond

The examples show that neither regularity, nor Theorem 3 remain true in general. It turns out that for some semantics we do not observe any differences if sticking to unrestricted AFs. In particular, this is the case for all considered semantics characterizable by the admissible or naive kernel. For the remaining semantics Theorem 3 becomes false but we may restore the characterization theorems if guaranteeing the existence of fresh arguments. We say that two frameworks F and G are *jointly expandable* (w.r.t. \mathcal{U}) if $\mathcal{U} \setminus (A(F) \cup A(G)) \neq \emptyset$.

Theorem 4 *For unrestricted AFs we have,*

1. $k(ad)$ characterizes \equiv_E^σ for any $\sigma \in \{ss, eg, ad, pr, il\}$,
2. $k(na)$ characterizes \equiv_E^σ for any $\sigma \in \{na, cf\}$.

For jointly expandable AFs we have,

3. $k(stb)$ characterizes \equiv_E^σ for any $\sigma \in \{stb, stg\}$ and
4. $k(\sigma)$ characterizes \equiv_E^σ for any $\sigma \in \{gr, co\}$.

The main proof strategies are straightforward extensions of those presented in [Oikarinen and Woltran, 2011]. However, one has to pay attention whenever properties of finite AFs like jointly expandability or finite computation of grounded semantics are used. Table 3 summarizes all results. The entry "k" in row *certain* AFs and column σ indicates that for *certain* AFs expansion equivalence w.r.t. semantics σ is characterizable through k . A question mark represents two different things. First, the characterization problem is unsolved and second, it can be checked that none of the introduced kernels serve as a characterization. As an aside, in consideration of Example 3 no node-preserving kernel may serve as a characterizing kernel for stable semantics in case of unrestricted AFs.

3.6 Realizability

Realizability is concerned with the expressive power of logical formalisms. It is well-known that in case of propositional logic any finite set of two-valued interpretations is realizable. Given such a finite set \mathcal{I} , we always find a set of formulae T , s.t. $Mod(T) = \mathcal{I}$. In abstract argumentation we are equipped

with a high number of semantics and it was shown that representational limits highly depend on the chosen semantics [Dunne *et al.*, 2015]. We will see that it is essential which class of frameworks we consider if we want to compare the expressiveness of two semantics. In case of finite AFs it does not make a difference whether we use semi-stable or preferred semantics but it does if sticking to unrestricted frameworks.

We proceed with the formal definition of different σ -signatures which collect all σ -realizable extension-sets of a certain class of AFs. Note that members of the compact versions require a witnessing framework without rejected arguments, i.e. any argument is contained in at least one extension.

Definition 4 The *unrestricted, finite, compact and finite compact* σ -signatures of a semantics σ are defined as follows:

1. $\Sigma_\sigma = \{\sigma(F) \mid F \text{ is an AF}\}$,
2. $\Sigma_\sigma^{\text{fin}} = \{\sigma(F) \mid F \text{ is a finite AF}\}$,
3. $\Sigma_\sigma^c = \{\sigma(F) \mid F \text{ is an AF}, A(F) = \bigcup \sigma(F)\}$ and
4. $\Sigma_\sigma^{c,\text{fin}} = \Sigma_\sigma^{\text{fin}} \cap \Sigma_\sigma^c$. \diamond

Clearly, for any semantics σ , $\Sigma_\sigma^{c,\text{fin}} \subseteq \Sigma_\sigma^c$, $\Sigma_\sigma^{\text{fin}} \subseteq \Sigma_\sigma$. For a detailed discussion on the relation between finite and finite compact signatures see [Baumann *et al.*, 2016a]. For some semantics both concepts coincide. Since finite AFs may realize finite extensions only, for any considered semantics σ a strict subset relation between finite compact and compact as well as finite and unrestricted signatures. A detailed analysis is part of future work. Let us recall a comparison of the expressiveness of the semantics in the realm of finiteness.

Theorem 5 [Dunne *et al.*, 2015] For *finite signatures*,

1. $\Sigma_{na}^{\text{fin}} \subset \Sigma_{stg}^{\text{fin}} = \Sigma_{stb}^{\text{fin}} \setminus \{\emptyset\} \subset \Sigma_{ss}^{\text{fin}} = \Sigma_{pr}^{\text{fin}}$,
2. $\Sigma_{cf}^{\text{fin}} \subset \Sigma_{ad}^{\text{fin}} \subset \Sigma_{co}^{\text{fin}}$ and
3. $\{\{E\} \mid E \subseteq \mathcal{U}, |E| \in \mathbb{N}\} = \Sigma_\sigma^{\text{fin}} \subset \Sigma_\tau^{\text{fin}}$ where $\sigma \in \{gr, il, eg\}$ and $\tau \in \{na, stb, stg, pr, ss, co\}$.

Sticking to unrestricted AFs makes a comparison more complex and for a long time it was unclear how to systematically achieve comparison results regarding expressiveness. It is the main result of this section that the relation between finite compact signatures is intimately connected to their relation in case of unrestricted signatures. More precisely, non-empty relative complements in case of finite compact signatures between two semantics carry over to their unrestricted versions.

Theorem 6 Given two semantics $\sigma, \tau \in \{na, stb, stg, ss, pr, co, gr, il, eg, cf, ad\}$ we have:

1. If $\Sigma_\sigma^{c,\text{fin}} \setminus \Sigma_\tau^{c,\text{fin}} \neq \emptyset$, then $\Sigma_\sigma^c \setminus \Sigma_\tau^c \neq \emptyset$ as well as
2. If $\Sigma_\sigma^c \setminus \Sigma_\tau^c \neq \emptyset$, then $\Sigma_\sigma \setminus \Sigma_\tau \neq \emptyset$.

Now we are prepared for a comparison in case of unrestricted frameworks.

Theorem 7 For *unrestricted signatures* the following hold:

1. $\{\{E\} \mid E \subseteq \mathcal{U}\} = \Sigma_\sigma \subset \Sigma_{na} \subset \Sigma_\tau$ for $\sigma \in \{gr, il\}$ and $\tau \in \{stb, stg, ss, pr\}$,
2. $\Sigma_{eg} \subset \Sigma_{pr}$,
3. $\Sigma_{stb} \subset \Sigma_\sigma$ for $\sigma \in \{stg, ss\}$,
4. $\Sigma_{pr} \setminus (\Sigma_{stb} \cup \Sigma_{ss} \cup \Sigma_{stg}) \neq \emptyset$,
5. $\Sigma_{stg} \setminus (\Sigma_{stb} \cup \Sigma_{pr} \cup \Sigma_{ss}) \neq \emptyset$,
6. $\Sigma_{stb} \setminus \Sigma_{pr} \neq \emptyset$,

7. $\Sigma_{ss} \setminus (\Sigma_{stb} \cup \Sigma_{pr} \cup \Sigma_{stg}) \neq \emptyset$,
8. $\Sigma_{co} \setminus \Sigma_\sigma \neq \emptyset$ and $\Sigma_\sigma \setminus \Sigma_{co} \neq \emptyset$ for $\sigma \in \{cf, ad\}$,
9. $\Sigma_{cf} \subset \Sigma_{ad}$.

In order to highlight the tremendous change we additionally present the relations in Venn-diagrams for some selected semantics. The dashed areas represent particular intersections for which the question of non-emptiness has to be left open.




Figure 2: Signatures of Selected Semantics.

3.7 Intertranslatability

Intertranslatability revolves around the idea of mapping one semantics to another. A motivation from [Dvořák and Woltran, 2011] is to reuse a solver for one semantics for another. The main tool for this endeavour are functions mapping AFs to AFs, so-called *translations* formally defined as follows.

Definition 5 [Dvořák and Woltran, 2011] Given two semantics σ, τ . A function $f : \mathcal{F} \rightarrow \mathcal{F}$ is called an *exact translation* for $\sigma \rightarrow \tau$, if $\sigma(F) = \tau(f(F))$ for any AF F . It is called a *faithful translation* if for any AF F first $|\sigma(F)| = |\tau(f(F))|$ and second $\sigma(F) = \{S \cap A(F) \mid S \in \tau(f(F))\}$. \diamond

Please note that for some semantics there are no exact translations available due to reasons inherent to those semantics. For instance, preferred semantics satisfies *I-maximality*, i.e. for any AF F , $pr(F)$ forms a \subseteq -antichain [Baroni and Giacomin, 2007]. This implies that an exact translation $ad \rightarrow pr$ can not exist since for $F = (\{a\}, \emptyset)$ we observe $\{\emptyset, \{a\}\} = ad(F)$. Sticking to faithful translations provides us with a positive answer if we consider finite AFs only [Spanring, 2012, Translation 3.1.85]. However, the considered translation does not serve in the general unrestricted case and interestingly, we will see that any search for a suitable translation will never succeed.

We proceed with some preliminary exact translation results.

Theorem 8 Given an *unrestricted AF* F . Define $Tr_{ad}^{cf}(F) = G$ with $A(G) = \{a \in A(F) \mid (a, a) \notin R(F)\}$ and $R(G) = \{(a, b), (b, a) \mid a, b \in A(G), (a, b) \in R(F)\}$. We have that $cf(F) = ad(G)$ and $na(F) = \sigma(F)$ for $\sigma \in \{pr, ss, stg, stb\}$.


According to Table 1 grounded and ideal semantics result in exactly one extension even in case of unrestricted AFs. Eager semantics in contrast additionally exhibits the possibility of infinitely many extensions. In this case eager and preferred semantics coincide [Baumann and Spanring, 2015, Theorem 15]. Both results give rise for the following exact translations.

Theorem 9 Given an *unrestricted AF* F and let $\sigma \in \{gr, il\}$, then $Tr_\tau^\sigma(F) = (S, \emptyset)$ is an exact translation for $\sigma \rightarrow \tau$ with $\tau \in \{co, pr, stg, ss, stb, gr, il, eg\}$ and $\sigma(F) = \{S\}$. For $eg \rightarrow pr$ use $Tr_\tau^{eg}(F)$ in case $|eg(F)| = 1$ and id otherwise.

Remember that stable semantics may collapse even for finite AFs. This means, there are frameworks which do not provide

any stable extension. In order to provide exact translations for $stb \rightarrow \tau$ we have to construct an AF G , s.t. $\tau(G) = \emptyset$.

Example 5 (Semantical Collapse) Consider the non-finitary AF $G = (A, R_1 \cup R_2 \cup R_3)$ as indicated below. Formally, $A = \{p_i, q_i, r_i \mid i \in \mathbb{N}\}$, $R_1 = \{(p_i, q_i), (q_i, p_i), (p_i, r_i) \mid i \in \mathbb{N}\}$, $R_2 = \{(r_k, r_{k+1}), (r_{k+1}, r_{k+2}), (r_{k+2}, r_k) \mid 0 \equiv k \pmod{3}\}$ and $R_3 = \{(p_i, p_j), (p_j, r_i) \mid i < j, i \equiv j \pmod{3}\}$.



Due to lack of space we omit a detailed semantical analysis. However, it can be checked that $A = \bigcup pr(F)$ and more importantly, $stb(F) = stg(F) = ss(F) = \emptyset$. Since $\bigcap \emptyset = \mathcal{U}$ we further infer $eg(F) = pr(F)$. \diamond

Theorem 10 Let $\sigma \in \{stg, ss, stb\}$, F an *unrestricted AF*. For G from Example 5 we have $\sigma(Tr_\sigma^{stb}(F)) = stb(F)$ with

$$Tr_\sigma^{stb}(F) = \begin{cases} F, & stb(F) \neq \emptyset; \\ G, & \text{otherwise.} \end{cases}$$

So far, we gave concrete translations only. The following theorem generalizes the finite version from [Dvořák and Spanring, 2016, Section 6.1]. It establishes a close relation between realizability and intertranslatability: if τ is not less expressive than σ , then σ can be exactly translated to τ and vice versa.

Theorem 11 Given semantics σ, τ , we have $\Sigma_\sigma \subseteq \Sigma_\tau$ if and only if there is an exact translation for $\sigma \rightarrow \tau$.

Note that this theorem implies that translational impossibilities entail differences in the expressive power. For instance, the following example shows that (in contrast to finite AFs) there is no faithful translation for $ad \rightarrow pr$ yielding $\Sigma_{ad} \setminus \Sigma_{pr} \neq \emptyset$.

Example 6 (Translational Impossibility) The AF $F = (\mathcal{U}, \emptyset)$ prohibits a faithful (and thus exact) translation Tr for $ad \rightarrow pr$. This is due to the fact that firstly, $\mathcal{U} \in ad(F)$ enforces $Tr(F) = F$ because of the lack of fresh arguments and secondly, $pr(F) = \{\mathcal{U}\}$ entails $Tr(F) \neq F$. \diamond

Instead of listing all translational (im)possibilities we provide figures illustrating all results in an eye-catching way. Figure 3b summarizes known results regarding faithful translations in the finite case [Dvořák and Woltran, 2011; Spanring, 2012; Dvořák and Spanring, 2016], augmented with obvious insights for unique status semantics il and eg . For semantics σ, τ , encirclement in the same component indicates bidirectional translations. An arrow from σ to τ means directional translations. If there is no directed path (for instance for na to cf , or for cf to gr), then there is no translation. Figure 3a features the same visualization for unrestricted AFs. Dropping the restriction assumption has some further consequences for the considered semantics, namely exact and faithful intertranslatability coincide. It is a part of future work to study

which basic requirements a semantics has to fulfill, s.t. both forms of translations are essentially the same. However, in consideration of Theorem 11 we may interpret Figure 3a as a comparison of the expressiveness of the considered semantics. That is, $\Sigma_\sigma \subseteq \Sigma_\tau$ if and only if there is a directed path from σ to τ . This interpretation supplements the former illustration depicted in Venn-diagram 2b.

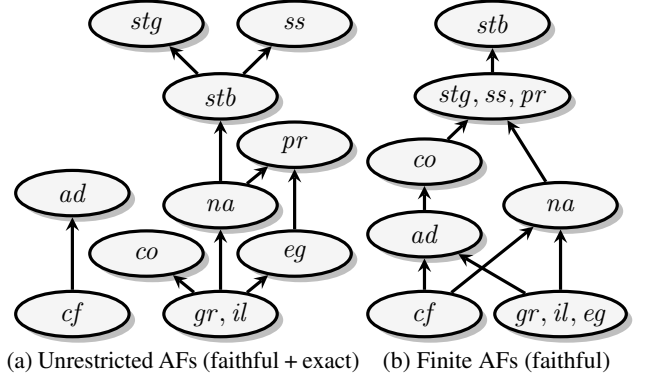


Figure 3: Intertranslatability.

4 Discussion

We provided a first extensive overview of abstract properties in case of unrestricted frameworks. In almost all cases we observed substantial differences if leaving the realm of finiteness. However, it is an important observation of this study that some differences can be overcome if considering AFs expandable by a certain number of “fresh” arguments. For instance, in case of expansion equivalence we have seen that being jointly expandable is sufficient for restoring former results (Theorem 4). Moreover, regarding translations, we observed that for slightly restricted AFs $F = (A, R)$, s.t. $|A| \leq |\mathcal{U} \setminus A|$ we are able to provide exact and efficiently computable translations from preferred to semi-stable semantics via $Tr_{ss}^{pr}(F) = F' = (A', R')$ with $A' = A \cup \{a' \mid a \in A\}$ and $R' = R \cup \{(a, a'), (a', a') \mid a \in A\}$. It is part of future work to figure out whether this restriction allows for similar translational possibilities as in case of finite AFs (cf. Figure 3b).

Another promising issue related to the presented research are alternative characterizations of $cf2$ and $stage2$ semantics as presented in [Gaggl and Woltran, 2013; Dvořák and Gaggl, 2012]. These definitions avoid the recursive computation of subframeworks and it is an interesting question whether such alternative characterizations may overcome problems of undefinedness as discussed in Section 3.4.

Finally, it is already known that AFs can be seen as a restricted class of logic programs (LPs). More precisely, there is a standard translation Tr from AFs to LPs, s.t. $\sigma(F)$ and $\tau(Tr(F))$ coincide for certain pairs of semantics σ and τ [Strass, 2013, Theorem 4.13]. Consequently, one interesting research question is to which extent results regarding infinite AFs may contribute to a better understanding of infinite LPs.

Acknowledgements

This research has been supported by the Austrian Science Fund (FWF) through project P25521.

References

- [Baroni and Giacomin, 2007] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, 171:675–700, 2007.
- [Baroni *et al.*, 2005] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. Scc-recursiveness: A general schema for argumentation semantics. *Artif. Intell.*, 168(1-2):162–210, 2005.
- [Baroni *et al.*, 2011] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *KER*, 26:365–410, 2011.
- [Baroni *et al.*, 2013] Pietro Baroni, Federico Cerutti, Paul E. Dunne, and Massimiliano Giacomin. Automata for infinite argumentation structures. *Artif. Intell.*, 203:104–150, 2013.
- [Baroni *et al.*, 2014] Pietro Baroni, Massimiliano Giacomin, and Beishui Liao. On topology-related properties of abstract argumentation semantics. A correction and extension to dynamics of argumentation systems: A division-based method. *Artif. Intell.*, 212:104–115, 2014.
- [Baumann and Spanring, 2015] Ringo Baumann and Christof Spanring. Infinite argumentation frameworks - on the existence and uniqueness of extensions. In *Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, volume 9060, pages 281–295. Springer, 2015.
- [Baumann and Strass, 2012] Ringo Baumann and Hannes Strass. Default reasoning about actions via abstract argumentation. In *COMMA*, pages 297–309, 2012.
- [Baumann and Strass, 2016] Ringo Baumann and Hannes Strass. An abstract logical approach to characterizing strong equivalence in logic-based knowledge representation formalisms. In *KR*, pages 525–528, 2016.
- [Baumann *et al.*, 2016a] Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, Christof Spanring, Hannes Strass, and Stefan Woltran. On rejected arguments and implicit conflicts: The hidden power of argumentation semantics. *Artif. Intell.*, 241:244–284, 2016.
- [Baumann *et al.*, 2016b] Ringo Baumann, Thomas Linsbichler, and Stefan Woltran. Verifiability of argumentation semantics. In *COMMA*, pages 83–94, 2016.
- [Baumann, 2011] Ringo Baumann. Splitting an argumentation framework. In *LPNMR*, volume 6645 of *Lecture Notes in Computer Science*, pages 40–53. Springer, 2011.
- [Baumann, 2012] Ringo Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In *ECAI*, volume 242, pages 127–132. IOS Press, 2012.
- [Belardinelli *et al.*, 2015] Francesco Belardinelli, Davide Grossi, and Nicolas Maudet. Formal analysis of dialogues on infinite argumentation frameworks. In *IJCAI*, pages 861–867, 2015.
- [Bench-Capon and Dunne, 2007] Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artif. Intell.*, 171(10-15):619–641, 2007.
- [Caminada and Oren, 2014] Martin W. A. Caminada and Nir Oren. Grounded semantics and infinitary argumentation frameworks. In *BNAIC*, 2014.
- [Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77:321–357, 1995.
- [Dunne *et al.*, 2015] Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.*, 228:153–178, 2015.
- [Dvořák and Spanring, 2016] Wolfgang Dvořák and Christof Spanring. Comparing the expressiveness of argumentation semantics. *JLogCom*, page exw008, 2016.
- [Dvořák and Woltran, 2011] Wolfgang Dvořák and Stefan Woltran. On the intertranslatability of argumentation semantics. *J. Artif. Intell. Res.*, 41:445–475, 2011.
- [Dvořák and Gaggl, 2012] Wolfgang Dvořák and Sarah Alice Gaggl. Incorporating stage semantics in the scc-recursive schema for argumentation semantics. In *NMR*, 2012.
- [Gabbay and Grossi, 2014] Dov M. Gabbay and Davide Grossi. When are two arguments the same? equivalence in abstract argumentation. In *Johan van Benthem on Logic and Information Dynamics*, pages 677–701. 2014.
- [Gaggl and Strass, 2014] Sarah Alice Gaggl and Hannes Strass. Decomposing Abstract Dialectical Frameworks. In *COMMA*, volume 266, pages 281–292. IOS Press, 2014.
- [Gaggl and Woltran, 2013] Sarah Alice Gaggl and Stefan Woltran. The cf2 argumentation semantics revisited. *JLogCom*, 23:925–949, 2013.
- [Oikarinen and Woltran, 2011] Emilia Oikarinen and Stefan Woltran. Characterizing strong equivalence for argumentation frameworks. *Artif. Intell.*, 175:1985–2009, 2011.
- [Prakken, 2010] Henry Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2010.
- [Rahwan and Simari, 2009] Iyad Rahwan and Guillermo R. Simari. *Argumentation in Artificial Intelligence*. Springer Publishing Company, Incorporated, 1st edition, 2009.
- [Spanring, 2012] Christof Spanring. Intertranslatability results for abstract argumentation semantics. Master’s thesis, uniwiien, 2012.
- [Strass, 2013] Hannes Strass. Approximating operators and semantics for abstract dialectical frameworks. *Artif. Intell.*, 205:39–70, December 2013.
- [Strass, 2015] Hannes Strass. Instantiating rule-based defeasible theories in abstract dialectical frameworks and beyond. *JLogCom*, 2015. In press.
- [Weydert, 2011] Emil Weydert. Semi-stable extensions for infinite frameworks. In *BNAIC*, pages 336–343, 2011.